

Whole-Volume Clustering of Time Series Data from Zebrafish Brain Calcium Images via Mixture Model-Based Functional Data Analysis

Hien D. Nguyen*, Jeremy F. P. Ullmann, Geoffrey J. McLachlan, Venkatakaushik Voleti, Wenze Li,
Elizabeth M. C. Hillman, David C. Reutens, and Andrew L. Janke

Abstract

An increasing number of techniques are enabling visualization of neurological activity in animal models, generating large volumes of spatially correlated time series. A model-based functional data analysis technique for the clustering of such data is proposed and justified. An example analysis of a zebrafish imaging experiment is presented.

Index Terms

Calcium imaging; Functional data analysis; Mixture model; Zebrafish

I. INTRODUCTION

Calcium is a ubiquitous secondary messenger that regulates vital signaling events in neurons and other excitable cells [1]. In neurons, synaptic input and action potential firing trigger the rapid input of large quantities of calcium ions [2]. By fusing a calcium-binding domain to a fluorescent protein it is possible to monitor calcium ions and therefore various spike features [3]. Genetically encoded calcium indicators (GECIs) can target specific cell types or subcellular compartments and have facilitated long-term optical recording of genetically targeted neurons in vivo. GECIs like GCaMP are the most widely used calcium indicators [4] and a range of mouse, drosophila, and zebrafish lines have been generated that express GCaMP pan-neuronally.

In recent years, the zebrafish has become a popular model in neurological research (e.g. [5] and [6]) and the number of studies imaging neuronal calcium dynamics in larval zebrafish has increased dramatically (e.g. [7] and [8]). Larval zebrafish are particularly amenable to calcium-imaging as they are translucent making it easy to image large populations of neurons without surgery, possess a relatively small brain ($\approx 800 \times 600 \times 200 \mu\text{m}^3$) thereby

Hien Nguyen and Geoffrey McLachlan are with the School of Mathematics and Physics, The University of Queensland, St. Lucia, Brisbane, Australia 4075. Hien Nguyen, David Reutens, and Andrew Janke are with the Centre for Advanced Imaging, The University of Queensland, St. Lucia, Brisbane, Australia 4075. Jeremy Ullmann is with the Department of Neurology, Boston Children's Hospital and Harvard Medical School, Boston, USA 02115. Venkatakaushik Voleti, Wenze Li, and Elizabeth Hillman are with the Laboratory for Functional Optical Imaging, Departments of Biomedical Engineering and Radiology, Columbia University, New York, New York 10027, USA. *Corresponding Author: Hien Nguyen (Email: h.nguyen7@uq.edu.au).

facilitating whole-brain imaging [9], and the generation of transgenic lines is relatively straightforward promoting the creation of a range of transgenic lines [10]. Light-sheet based microscopy methods can enable imaging in larval zebrafish with sufficient temporal resolution to record neural dynamics throughout the brain [5].

In this article, we devise a method for clustering time series that are obtained from whole-volume calcium imaging experiments. Large time series data sets, consisting of thousands of time points, were acquired using swept confocally-aligned planar excitation microscopy (SCAPE) [11]. SCAPE is a novel method for high-speed light sheet imaging which enables imaging of the entire zebrafish brain using a single objective lens in a flexible geometry and at very high spatiotemporal resolution. Data sets utilized here were whole-volume calcium images of the of the larval zebrafish brain at a spatial resolution of $640 \times 130 \times 300$ ($x \times y \times z$) voxels and a temporal resolution of 4 volumes per second.

Two primary factors make the clustering of whole-volume calcium imaging time series difficult. Firstly, the time series can often be obtained at sparse or irregular intervals. Secondly, the size of the data firmly places the problem within the realms of Big Data, where conventional methods are infeasible to implement (cf. [12, Sec. 1.2]). The method that we develop in this article is designed to address these primary concerns.

Time series clustering is a well-studied area in data analysis and signals processing; see [13] and [14] for recent literature reviews. From [14], it is clear that there are numerous approaches to the problem of time series clustering. Due to the high dimensionality of the raw time series that arise from calcium imaging experiments, an abstraction is necessary to limit computational burden. In this article, we suggest the general framework of functional data analysis (FDA); see [15]. In this article, we present a mixture model-based functional data approach; thus only a brief literature review of such methods will be provided in the sequel.

In [16], a mixture of polynomial basis-filtered models was suggested for the clustering of time-dependent trajectories. In [17], a regularized mixture of mixed-effects (MEM) basis-filtered models using B-splines was considered for the clustering of sparsely-sampled univariate functional data. Both [18] and [19] suggested MEM basis-filtered models using Fourier bases for clustering of time-course gene expression data; [20] extended the method [19] by allowing for autoregressive (AR) errors. An extension of the MEM model of [19] that allowed for concomitant covariates was proposed in [21]. An adaptation of the MEM model from [17] was suggested for the clustering of sparsely-sampled bivariate functional data in [22]. Lastly, a set of methods that are related to the mixture model-based FDA approach are the mixtures of AR models that utilize a parametric stochastic process abstraction of the time series instead; see [23], [24], and [25].

In this article, we consider a two-stage approach to clustering the calcium image obtained time series that is adapted from that of [26]. In the first stage, each time series is filtered via a common set of B-spline bases using ordinary least squares (OLS); this produces a set of basis coefficients for each time series. In the second stage, a GMM-like (Gaussian mixture model [27, Ch. 3]) clustering of the basis coefficients for each time series is conducted to partition the time series into k classes, where k is determined by a BIC-like (Bayesian information criterion [28]) approach via the slope heuristic technique of [29]. In this article, we use the trimmed k -means algorithm [30] for clustering due to its speed in the large data context. The two-stage nature of our method shares similarities with the methods of [31] and [32].

To justify our two-stage approach, we demonstrate that it permits the same statistical interpretation as the MEM model of [17]. Although the method of [17] permits sparseness and irregularity in the time series, the iterative and slow-converging EM (expectation–maximization [33]) algorithm that is required for its implementation is inappropriate for data sizes in calcium imaging time series clustering; the same reasoning applies to the infeasibility of the MEM model-based methods; see [26, Sec. 4] for details. The mixture of AR model-based methods are also inappropriate in this context due to their inability to handle sparsity in the sampling of the time series.

Along with the statistical model for calcium imaging time series clustering and two-stage method for performing the clustering, we also discuss the theoretical conditions under which the maximum likelihood (ML) estimator of a GMM model is consistent, under dependence between spatially correlated voxels. We then extend this discussion towards the suitability of the trimmed k -means algorithm. To demonstrate our methodology, we conduct a clustering of time series data arising from a whole-volume calcium imaging experiment of a larval zebrafish brain at rest, acquired using SCAPE microscopy.

The rest of the article proceeds as follows. A statistical model for calcium imaging time series clustering is presented in Section II. A two-stage algorithm is suggested for performing clustering in Section III. An application to a zebrafish brain calcium imaging data is presented in Section IV. Conclusions are drawn in Section V. A description of the trimmed k -means method is also provided in the Appendix.

II. STATISTICAL MODEL

Suppose that we wish to cluster a sample of time series $\mathcal{Y}_n = \{Y_1(t), \dots, Y_n(t)\}$ that are treated as functions in t on some interval \mathbb{T} . Let each time series $Y_i(t)$ ($i = 1, \dots, n$) be observed at m_i time points $t_1, \dots, t_{m_i} \in \mathbb{T}$ and suppose that we observe Y_i at t_j ($j = 1, \dots, m_i$) indirectly via the noise-corrupted variable

$$Z_i(t_j) = Y_i(t_j) + E_i(t_j), \quad (1)$$

where $E_i(t_j)$ is a random error from a univariate Gaussian density distribution with mean 0 and variance σ^2 .

In order to cluster, we require an implicit model for the subpopulation heterogeneity of the elements in \mathcal{Y}_n . We propose a parametric model for the purpose. Let $Y_i(t) = \mathbf{B}_i^T \mathbf{x}(t)$, where $\mathbf{x}(t)$ is a d -dimensional vector of evaluations at point t of a B-spline system with $(d-2)$ -breakpoints over \mathbb{T} (see [34, Ch. 9]), and let $\mathbf{B}_i \in \mathbb{R}^d$ be a random variable that captures the heterogeneity of \mathcal{Y}_n . Here, the superscript T indicates matrix transposition. We suppose that the B-spline representation \mathbf{B}_i arises from a k -component GMM with density function

$$f(\mathbf{b}; \boldsymbol{\theta}) = \sum_{c=1}^k \pi_c \phi_d(\mathbf{b}; \boldsymbol{\mu}_c, \mathbf{V}_c), \quad (2)$$

where $\pi_c > 0$, $\sum_{c=1}^k \pi_c = 1$, $\boldsymbol{\mu}_c \in \mathbb{R}^d$, $\mathbf{V}_c \in \mathbb{R}^{d \times d}$ is positive definite ($c = 1, \dots, k$), and

$$\phi_d(\mathbf{b}; \boldsymbol{\mu}, \mathbf{V}) = |2\pi\mathbf{V}|^{-1/2} \exp \left[-\frac{1}{2} (\mathbf{b} - \boldsymbol{\mu})^T \mathbf{V}^{-1} (\mathbf{b} - \boldsymbol{\mu}) \right]$$

is the multivariate Gaussian density function with mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{V} . We shall refer to π_c as the prior probability of observing Y_i is in cluster c , and $\phi_d(\mathbf{b}; \boldsymbol{\mu}, \mathbf{V})$ as the component density function of cluster c . Here, we put the parameter components π_c , $\boldsymbol{\mu}_c$, and \mathbf{V}_c into the vector $\boldsymbol{\theta}$.

Remark 1. Any linear basis system can be used in place of the B-splines in this article. Examples of other basis systems are Fourier or polynomial bases; see [15, Sec. 3.3].

The model described so far is close to that of [17]. However, in order to proceed [17] estimation of each \mathbf{B}_i is required using the data $Z_i(t_j)$, for all i and j , via an EM algorithm. This is not feasible in the setting that we wish to conduct clustering. We instead suggest that we estimate each \mathbf{B}_i only via the data $Z_i(t_j)$ for the same i . We do this via the OLS estimator

$$\tilde{\mathbf{B}}_i = (\mathbf{X}_i^T \mathbf{X}_i)^{-1} \mathbf{X}_i^T \mathbf{Z}_i, \quad (3)$$

where we let $\mathbf{X}_i^T = \begin{bmatrix} \mathbf{x}(t_1) & \cdots & \mathbf{x}(t_{m_i}) \end{bmatrix}$ and $\mathbf{Z}_i^T = (Z_i(t_1), \dots, Z_i(t_{m_i}))$. The following result is proved in [26].

Proposition 2. *Under characterizations (1) and (2), $\hat{\mathbf{B}}_i$ [as characterized by (3)] has density function*

$$f(\tilde{\mathbf{b}}_i; \boldsymbol{\vartheta}) = \sum_{c=1}^k \pi_c \phi_d(\tilde{\mathbf{b}}_i; \boldsymbol{\mu}_c, \mathbf{V}_c + \sigma^2 [\mathbf{X}_i^T \mathbf{X}_i]^{-1}), \quad (4)$$

where we put the parameter components π_c , $\boldsymbol{\mu}_c$, \mathbf{V}_c , and σ^2 into $\boldsymbol{\vartheta}$.

Remark 3. Density (4) is dependent on i only via the term $\sigma^2 [\mathbf{X}_i^T \mathbf{X}_i]^{-1}$. If we let $\mathbf{X}_i = \mathbf{X}$ for all i , then we can write

$$f(\tilde{\mathbf{b}}; \boldsymbol{\psi}) = \sum_{c=1}^k \pi_c \phi_d(\tilde{\mathbf{b}}; \boldsymbol{\mu}_c, \tilde{\mathbf{V}}_c), \quad (5)$$

where $\tilde{\mathbf{V}}_c = \mathbf{V}_c + \sigma^2 [\mathbf{X}^T \mathbf{X}]^{-1}$ and $\boldsymbol{\psi}$ contains the parameter components of the model.

A. Parameter Estimation

Suppose that we observe data \mathbf{z}_i that are realizations of \mathbf{Z}_i from the process (1). Also suppose that $\{t_1, \dots, t_{m_i}\} = \{t_1, \dots, t_m\}$ for all i (i.e. $\mathbf{X}_i = \mathbf{X}$ for all i , from Remark 2). A natural parameter estimator for $\boldsymbol{\psi}$ in (5) is the MLE (maximum likelihood estimator). Define ∇ as the gradient operator, and let the MLE be $\hat{\boldsymbol{\psi}}_n$, which is a suitable root of the likelihood $\ell_n(\boldsymbol{\psi}) = \sum_{i=1}^n \log f(\tilde{\mathbf{b}}_i; \boldsymbol{\psi})$ formed on the basis of the data $\tilde{\mathbf{b}}_i = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{z}_i$. The following useful result regarding the consistency of MLE under dependence (e.g. spatial dependence between voxels in imaging data) is obtained in [26].

Proposition 4. *Let $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ be an identically distributed (ID) and stationary ergodic (or strong-mixing) random sample [as characterized by (1)], which generates a set of OLS estimators $\tilde{\mathbf{B}}_1, \dots, \tilde{\mathbf{B}}_n$. Suppose that each $\tilde{\mathbf{B}}_i$ ($i = 1, \dots, n$) arises from a distribution with density $f(\tilde{\mathbf{b}}; \boldsymbol{\psi}_0)$, where $\boldsymbol{\psi}_0$ (containing π_{0c} , $\boldsymbol{\mu}_{0c}$, and \mathbf{V}_{0c} for all c) is a strict-local maximizer of $\mathbb{E} \log f(\tilde{\mathbf{b}}; \boldsymbol{\psi})$. If $\Psi_n = \{\boldsymbol{\psi} : \nabla \ell_n = \mathbf{0}\}$ (where we take $\Psi_n = \{\bar{\boldsymbol{\psi}}\}$, for some $\bar{\boldsymbol{\psi}}$ in the domain of $f(\tilde{\mathbf{b}}; \boldsymbol{\psi})$, if $\nabla \ell_n = \mathbf{0}$ has no solution), then for any $\epsilon > 0$, $\lim_{n \rightarrow \infty} \mathbb{P} \left[\inf_{\boldsymbol{\psi} \in \Psi_n} (\boldsymbol{\psi} - \boldsymbol{\psi}_0)^T (\boldsymbol{\psi} - \boldsymbol{\psi}_0) > \epsilon \right] = 0$.*

The result establishes the fact that there exists a consistent root to the likelihood score equation $\nabla \ell_n = \mathbf{0}$. This is a useful result since the likelihood for a GMM cannot possess a unique global maximum as it is unbounded

and lacks identifiability. The proof of Proposition 4 invokes an extremum estimator theorem [35, Thm. 4.1.2] in conjunction with either an ergodic continuous mapping theorem or a strong-mixing continuous mapping theorem (see [36, Thms. 3.35 and 3.49]). A generic uniform law of large numbers (ULLN) such as [37, Thm. 5] can then be used to obtain the desired result; see [24, Appendix I] for the proof of a similar result.

Remark 5. Strong mixing can be guaranteed by assuming that the sequence \mathbf{Z}_i is M -dependent (cf. [38, Sec. 2.1]). That is, there exists some $M < \infty$ such that if $|i - i'| > M$ then \mathbf{Z}_i and $\mathbf{Z}_{i'}$ are independent. This is a reasonable assumption in imaging applications where features behave in locally coherent groups.

Remark 6. The intended use of this article is for clustering time series at voxels on a three-dimensional array. The asymptotics of Proposition 4 in conjunction with Remark 5 applies when one of the three dimensions are extended infinitely. In practice, this is unimportant, but a theoretical labelling of the voxels that would allow the M -dependent assumption to hold when all three dimensions are extended would be difficult. One could instead treat the problem as an estimation problem over data arising from a stationary three-dimensional mixing random field. In such a case, the M -dependent assumption of Remark 5 can still be made (cf. [39]). The generic ULLN of [40] can be used in place of [37, Thm. 4] to establish an equivalent result.

B. Clustering via the GMM

Under (5) and following the approach from [27, Secs. 1.15.1 and 1.15.2], we have that

$$a_i = \arg \max_{c=1,\dots,k} \pi_{0c} \phi_d(\tilde{\mathbf{b}}_i; \boldsymbol{\mu}_{0c}, \tilde{\mathbf{V}}_{0c}) / f(\tilde{\mathbf{b}}_i; \boldsymbol{\psi}_0) \quad (6)$$

is an outright assignment of the i th observation according to the optimal (Bayes) rule of allocation. Since $\boldsymbol{\psi}_0$ is unknown, we can substitute $\hat{\boldsymbol{\psi}}_n$ in (6) to obtain a plug-in allocation rule.

Remark 7. Extending upon Remark 3, if we also make the assumption $\pi_c = 1/k$ and $\tilde{\mathbf{V}}_c = \lambda \mathbf{I}_d$ for all c and some $\lambda > 0$, then the clustering obtained from a k -means algorithm (e.g. [41]) approximates a clustering obtained via a GMM. The invariance of the clustering rule to the scaling parameter λ can be seen by inspecting the discriminant function between any two classes [42, Eqn. 3.3.7]. The similarities between the algorithms for conducting GMM and k -means clustering are highlighted in [43, Secs. 16.1.1 and 16.1.2].

C. Model Selection

Thus far, we have not commented on the selection of the number of components k as it cannot be conducted within the likelihood approach of Section II-A. An external wrapper-type method for selecting k is the penalized-contrast approach of [44]. In the context of this article, the method can be presented as follows.

Let $\mathbb{K} = \{k_1, \dots, k_K\}$ be a set of K possible values for the number of components and let $\hat{\boldsymbol{\psi}}_n^{[k]}$ be the MLE of (5) with $k \in \mathbb{K}$ components. Let $k_0 \in \mathbb{K}$ be the optimal number of components in the sense that it minimizes the expectation of the loss $n^{-1} \ell_n(\boldsymbol{\psi}_0^{[k_0]}) - n^{-1} \ell_n(\hat{\boldsymbol{\psi}}_n^{[k]})$ for $k \in \mathbb{K}$. We can estimate k_0 by

$$\hat{k} = \arg \min_{k \in \mathbb{K}} -\frac{1}{n} \ell_n(\hat{\boldsymbol{\psi}}_n^{[k]}) + 2\tilde{\kappa} \text{pen}(k), \quad (7)$$

where $\text{pen}(k)$ is a problem-specific penalization function for a model with k components and $\tilde{\kappa}$ is the so-called slope heuristic that is estimated from the data (cf. [45]). According to [46, Tab. 1], an appropriate penalty function for the GMM is to use the number of parameters. That is, under the restricts of Remark 3, $\text{pen}(k) = (d^2/2 + 3d/2 + 1)k - 1$. Furthermore, under the additional restrictions of Remark 7, we get $\text{pen}(k) = dk$. The slope heuristic $\tilde{\kappa}$ can be estimated via the DDSE (data-driven slope estimation) method of [46]; see also [29].

III. TWO-STAGE ALGORITHM

Make the assumptions from Remark 3. Let $\mathbf{z}_1, \dots, \mathbf{z}_n$ be a realization of the random sample of time series $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ that are observed at the n voxels, where $\mathbf{Z}_i^T = (Z_i(t_1), \dots, Z_i(t_m))$ and $t_j \in \mathbb{T}$ ($i = 1, \dots, n$ and $j = 1, \dots, m$). The first stage of the algorithm is to filter each \mathbf{Z}_i via a B-spline system. The second stage of the algorithm is to cluster the OLS estimators that are obtained upon B-spline filtering.

A. Stage 1: B-spline Filter

As in Section II, let $\mathbf{x}(t)$ be a d -dimensional vector of evaluates at point t of a B-spline system with $(d-2)$ -breakpoints over \mathbb{T} . In this article, we find that $d = 200$ is sufficient for our application. Further, let $\mathbf{X}^T = \begin{bmatrix} \mathbf{x}(t_1) & \cdots & \mathbf{x}(t_m) \end{bmatrix}$ be the matrix of B-spline system evaluates at the m sampled time points. For each i , we filter the time series \mathbf{z}_i via the OLS estimator $\tilde{\mathbf{b}}_i = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{z}_i$.

B. Stage 2: Clustering

Upon obtaining the OLS estimates, we now proceed to cluster the sample $\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n$. Unfortunately, even with a modest-size whole-volume brain calcium imaging dataset (e.g. $n \approx 10^6$ and $m \approx 2000$), estimating the MLE of (5) for use with Rule (6) is prohibitively computationally intensive for sufficiently large k (e.g. $k \geq 10$). Even the additional restrictions imposed in Remark 7 and the use of efficient k -means algorithm implementations do not improve the computational speed to an acceptable level.

Making the assumptions from Remark 7, we can sufficiently-quickly estimate the component means $\boldsymbol{\mu}_c$ for $c = 1, \dots, k$ (we will call the estimates $\tilde{\boldsymbol{\mu}}_{nc}$), for relatively large k (e.g. $k \geq 20$) via the trimmed k -means method of [30] as implemented via the TCLUST algorithm of [47]. A description of the trimmed k -means approach is provided in the Appendix.

Upon obtaining the estimates of all the component means $\tilde{\boldsymbol{\gamma}}_n^T = (\tilde{\boldsymbol{\mu}}_{n1}, \dots, \tilde{\boldsymbol{\mu}}_{nk})$, we can use Rule (6) to allocate each OLS estimate $\tilde{\mathbf{b}}_i$. Under the assumptions from Remark 7 and upon substitution of the component means, Rule (6) simplifies to the usual k -means rule

$$a_i = \arg \min_{c=1, \dots, k} \left(\tilde{\mathbf{b}}_i - \tilde{\boldsymbol{\mu}}_{nc} \right)^T \left(\tilde{\mathbf{b}}_i - \tilde{\boldsymbol{\mu}}_{nc} \right). \quad (8)$$

The value of k is determined by Rule (7) upon estimation of $\tilde{\kappa}$ via a sufficiently large sample from $\mathbb{K} = \{2, 3, \dots\}$. Here, the log-likelihood $\ell_n(\boldsymbol{\psi})$ (from Section II-A) simplifies to

$$\ell_n(\boldsymbol{\gamma}) = \sum_{i=1}^n \log \sum_{c=1}^k k^{-1} \phi_d \left(\tilde{\mathbf{b}}_i; \boldsymbol{\mu}_c, \lambda \mathbf{I}_d \right) \quad (9)$$

for some fixed $\lambda > 0$. In this article, we set $\lambda = 1$.




Figure 1. Time-averaged image of the 75th z -slice of the zebrafish brain calcium imaging volume acquired using SCAPE microscopy. The legend on the right-hand side indicates the average level of signal intensity at any particular voxel.

IV. EXAMPLE APPLICATION

A. Data Description

We consider an analysis of a time series data set arising from the volumetric calcium imaging of a larval zebrafish brain. The in-vivo calcium imaging was performed on a 5 day post fertilization *Tg(elavl3:H2B-GCaMP6s)* fish [48]. Images were acquired using SCAPE microscopy with 488 nm excitation [11]. The zebrafish was imaged at rest over approximately 20 minutes and $640 \times 130 \times 300$ ($x \times y \times z$) voxel volume (actual dimension $\approx 800 \times 600 \times 200 \mu\text{m}^3$) of time series data were acquired at 4 volumes per second.

As an example, we analyze a central region of interest within the brain, with volume $170 \times 70 \times 150$ ($x \times y \times z$) voxels. Each time point of each time series measures image intensity on a scale between 0 and 1. The total number of time points at each voxel of the data is $m = 1935$; due to irregularities, each series covers the wider time range of $t_1 = 1$ to $t_m = 2797$ after cropping. The volume contains a total of $n = 1785000$ spatially correlated time series from voxels potentially displaying interesting neuronal activity or anatomical features. A time averaged image of the 75th z -slice of the image is presented in Figure 1. Upon inspection of Figure 1, we see that there exist spatial variation in the mean signal that may indicate the presence of voxel subpopulations.

B. Data Analysis via Two-Stage Approach

All statistical analyses in this article are conducted within the *R* statistical environment [49] via the *Microsoft Open* 3.2.3 build. Data analysis is conducted using a mix of open source packages along with bespoke scripts. All computations are performed on a *MacBook Pro* (retina, 15-inch, early 2013) with a 2.4 GHz *Intel Core i7* processor and 16 GB of DDR3 RAM. Computation times are measured using the *proc.time* function in the *base* package of *R*.

Imaging data were first converted from the raw time frames (*ANDOR* format) using a custom *python* script. Part of this pre-processing involves a conversion to *MINC* [50], a HDF5 based format such that we can more easily




Figure 2. Panel A visualizes 10 randomly sampled time series of intensities from the total of $n = 1785000$ voxels; different colors indicate the signals from different time series. Panel B visualizes the $d = 100$ basis (98-node) B-spline representations $\tilde{Y}(t) = \tilde{\mathbf{b}}_i^T \mathbf{x}(t)$ for each of the 10 randomly sampled series; the colors correspond to those from Panel A.

deal with datasets that approximate 1TB. Once in the *MINC* format, the data is then realigned within each volume, realigned for each time series, and intensity normalized using the *MINC* tools. This post-processing is performed on a ≈ 200 core *Linux* cluster using *gridengine*. A sub-section of the data are then extracted and converted to *NIfTI* format for statistical analysis. The *NIfTI* format images were read into *R* via the *AnalyzeFMRI* package [51]. Before initializing Stage 1 of our process, we firstly detrend the data. This is performed rapidly via the *speedlm.fit* function within the *speedglm* package [52]. Using a $d = 100$ bases (i.e. a 98-node) B-spline system, the OLS estimates of the B-spline representation $\tilde{\mathbf{b}}_i$ ($i = 1, \dots, 1785000$) for Stage 1 of the procedure are quickly computed via the functions within the *fda* package [53]. An example sample of 10 time series $\mathbf{Z}_i^T = (Z_i(1), \dots, Z_i(2797))$ along with their estimated B-spline representations $\tilde{Y}(t) = \tilde{\mathbf{b}}_i^T \mathbf{x}(t)$ ($t = 1, \dots, 2797$) is visualized in Figure 2. Reading, detrending, and OLS estimation of every voxel in the volume was performed in a total computation time of 0.89 hours.

Prior to conducting Stage 2 of the process, we firstly column normalize the obtained OLS estimates $\tilde{\mathbf{b}}_i^T = (\tilde{b}_{i1}, \dots, \tilde{b}_{id})$. That is, for each $j = 1, \dots, d$, we normalize each of the j th column of coefficients $\tilde{b}_{1j}, \dots, \tilde{b}_{nj}$ by the respective mean and standard deviation. This is done to reduce the effects of differing dimensional scales. Upon

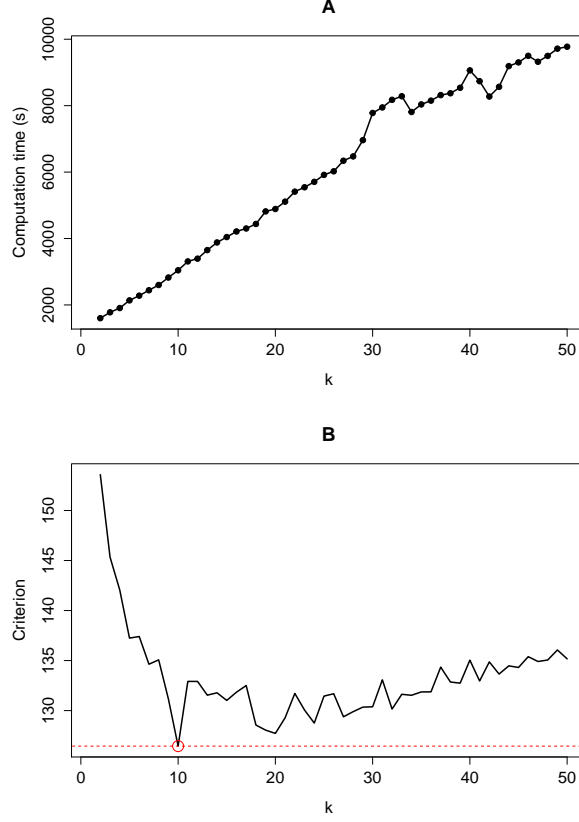


Figure 3. Panel A visualizes the total amount of time (in seconds) taken to run the TCLUST algorithm for each $k \in \mathbb{K}$ ($\mathbb{K} = \{2, \dots, 50\}$). Panel B visualizes the model selection criterion values for each k , as per Rule (10); the dotted line and circled point indicate the optimal value of the criterion and $\hat{k} = 10$, respectively.

normalization, the trimmed k -means algorithm is used to cluster the data. The TCLUST algorithm was applied via the *tkmeans* function from the *tclust* package [54]. Numbers of clusters in the set $\mathbb{K} = \{2, \dots, 50\}$ were considered, and the algorithm was repeated 20 times for each $k \in \mathbb{K}$, in order to mitigate against convergence to a spurious solution. The solution that maximized the objective (9) over all repetitions is taken to be the optimal k -cluster trimmed k -means estimator $\tilde{\gamma}_n^{[k]}$.

Using the objective sequence $\ell_n(\tilde{\gamma}_n^{[k]})$ ($k \in \mathbb{K}$), we utilize the *capushe* package [55] to implement the DDSE method of [29]. Under the penalty $\text{pen}(k) = dk$, The slope heuristic is estimated to be $\tilde{\kappa} = 1.335 \times 10^{-3}$. This results in the model selection rule

$$\hat{k} = \arg \min_{k \in \mathbb{K}} - \frac{\sum_{i=1}^n \log \sum_{c=1}^k \pi_c \phi_d(\tilde{\mathbf{b}}_i; \tilde{\boldsymbol{\mu}}_{cn}^{[k]}, \mathbf{I}_d)}{n} + \frac{2.67kd}{10^3}. \quad (10)$$

Under Rule 10, the optimal number of clusters is determined to be $\hat{k} = 10$. Panel A of Figure 3 visualizes the average time taken to perform 20 repetitions of the TCLUST algorithm for each $k \in \mathbb{K}$; Panel B of Figure 3 visualizes the model selection criterion values $-n^{-1} \ell_n(\tilde{\gamma}_n^{[k]}) + 2\tilde{\kappa} \text{pen}(k)$ for each considered k . The total computation time for 20 repetitions of the TCLUST algorithm for each k is 82.61 hours.




Figure 4. Visualization of the $\hat{k} = 10$ mean functions $\tilde{\mu}_c(t) = \tilde{\mu}_{cn}^T \mathbf{x}(t)$ ($c = 1, \dots, 10$) that are obtained via the TCLUST algorithm for trimmed k -means; the colors are as in Figure 5.

C. Results

Upon reversing the normalization discussed in Section IV-B, we can write the $\hat{k} = 10$ cluster mean functions as $\tilde{\mu}_c(t) = \tilde{\mu}_{cn}^T \mathbf{x}(t)$ ($c = 1, \dots, 10$). A plot of the 10 cluster mean functions appears in Panel A of Figure 4; a frequency plot of the number of voxels that are allocated into each cluster (via Rule (8)) appears in Panel B of Figure 4. From Figure 4, we notice that although the majority of mean functions appear parallel, there are a few that exhibit behavior that differs over time, and not simply differing in mean; for example, the $c = 3$ and $c = 9$ mean functions exhibit behavior that is substantially different to the others.

A visualization of the clustering at the 75th slice is presented in Figure 5. From Figure 5, we can make some inferences regarding the nature of the voxels that are allocated to each cluster. For example, $c = 1$ can be inferred as background or non-brain matter. Clusters $c = 2, 3, 4, 5, 6, 7$ can be inferred as edge effects on the interface between the background and the brain matter; the differences in these background effects may be explained by the various types of cellular materials such as membranes and tissues. Clusters $c = 8, 9, 10$ can be inferred as various types of brain matter. Furthermore, the cluster allocations in Figure 5 appears to be spatially coherent, thus indicating that the methodology is producing biologically meaningful results and not spurious allocations. Further, we observe that the frequencies of the different clusters are quite varied, indicating that the process is able to identify both rare and common subpopulations of voxels.

V. CONCLUSIONS

High-speed volumetric microscopy technologies hold great potential for the investigation of neurological activity in animals; see for example [11]. The methodology for analysis of whole-volume calcium imaging data has not been well developed and especially not adapted to the Big Data pathologies that are inherit in the class of data. In this article, we have developed a model-based clustering method that addresses the problems of sparse sampling, high dimensionality, and spatial dependency, which are inherent in whole-volume calcium imaging data. Our methodology is a two-stage mixture model-based approach for the functional data analysis of time series that arise from calcium imaging experiments.




Figure 5. Cluster allocations (as per Rule (8)) of the 75th z -slice of the zebrafish brain calcium imaging volume. The legend on the right-hand side indicates the cluster allocation and corresponds with the colors that are used in Figure 4.

We derived a theoretical model under which our approach can be interpreted. The approach is demonstrated to coincide with the estimation of a GMM for OLS estimates of B-spline coefficients. The consistency (under data dependence) of the MLE of the model parameter that is required for the application of the approach is provided. For feasibility, a simplification of the GMM is estimated using a trimmed k -means algorithm. A consistency result (under data dependence) is proven for the trimmed k -means estimators. This consistency result is novel and extends the previous results of [30] and [47] regarding the consistency of such estimators under independence of observations.

An example application of this method on a resting-state zebrafish calcium imaging dataset is presented. The computational timing of the approach, using freely-available open software, demonstrated the feasibility of the procedure. The outcome of the data analysis showed that the approach produces spatially coherent cluster allocations that are biologically interpretable and meaningful.

Future directions for this research will involve the application of the approach on animals undergoing stimulation; such data will likely exhibit more functional asynchronicities than those observed in the provided example. Computational times may be made faster via bespoke software for the problem, that are implemented in C, this may be incorporated in the current workflow via the *Rcpp* package [56]. The bespoke software may also make possible the application of less restrictive models for clustering.

Finally, although we note the spatial coherency of the obtained cluster allocations, such allocations can be made more spatially consistent via the application of an MRF (Markov random field) model, such as those successfully implemented in [57] and [24]. The MRF model can be applied as a third stage to the approach.

ACKNOWLEDGMENTS

This work was supported by a University of Queensland Early Career Award to JFPU, a NIH BRAIN initiative grant U01 NS094296 and DoD MURI W911NF-12-1-0594 to EH. HDN, GJM, and AJ are supported under the Australian Research Council’s Linkage Projects funding scheme (LP130100881). We would also like to thank

Misha Ahrens for providing us with zebrafish and Sharon Kim (Hillman lab) for assistance with fish protocols and organization.

APPENDIX

Description of Trimmed k -means

Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be a realization of an ID random sample $\mathbf{U}_1, \dots, \mathbf{U}_n \in \mathbb{R}^d$. Based on [47, Prop. 1], we can describe the α -trimmed k -means problem as the estimation of the mean vectors $\boldsymbol{\gamma}^T = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k)$, where $\boldsymbol{\mu}_c$ ($c = 1, \dots, k$), via the maximization of the objective

$$\mathcal{L}_n(\boldsymbol{\gamma}) = \frac{1}{n} \sum_{i=1}^n \sum_{c=1}^k \delta_c(\mathbf{u}_i; \boldsymbol{\gamma}) \log D_c(\mathbf{u}_i; \boldsymbol{\gamma}), \quad (11)$$

where $D_c(\mathbf{u}; \boldsymbol{\gamma}) = k^{-1} \phi_d(\mathbf{u}; \boldsymbol{\mu}_c, \lambda \mathbf{I}_d)$ and

$$\begin{aligned} \delta_c(\mathbf{u}; \boldsymbol{\gamma}) &= \mathbb{I} \left\{ \mathbf{u} \in \left\{ \max_{c'} D_{c'}(\mathbf{u}; \boldsymbol{\gamma}) = D_c(\mathbf{u}; \boldsymbol{\gamma}) \right\} \right. \\ &\quad \left. \cap \{D_c(\mathbf{u}; \boldsymbol{\gamma}) \geq R(\boldsymbol{\gamma})\} \right\}. \end{aligned} \quad (12)$$

Here $\lambda > 0$ is a fixed constant, $\mathbb{I}(A)$ equals 1 if A is true and 0 otherwise, and $R(\boldsymbol{\gamma}) = \inf_v \{G(v; \boldsymbol{\gamma}) \geq \alpha\}$, where $G(v; \boldsymbol{\gamma}) = \mathbb{P}(\max_c D_c(\mathbf{u}; \boldsymbol{\gamma}) \leq v)$ under the distribution of \mathbf{U}_1 . Call the α -trimmed k -means estimate $\tilde{\boldsymbol{\gamma}}_n^T = (\tilde{\boldsymbol{\mu}}_{n1}, \dots, \tilde{\boldsymbol{\mu}}_{nk})$.

Upon inspecting the objective (11), we observe that (12) sets the contribution of all observations \mathbf{u}_i that have $\max_c D_c(\mathbf{u}_i; \boldsymbol{\gamma})$ values below the α -quantile to 0; traditionally, the observations below the α -quantile are seen to be outliers. Thus, only $(1 - \alpha) \times 100\%$ of the data contributes non-zero values to the computation of the objective. This explains the increased computational speed of the trimmed k -means over its untrimmed counterpart, when α is sufficiently large. With respect to the size of data in our application, we found $\alpha = 0.9$ to be reasonable.

TCLUST Algorithm

Like the k -means problem, the α -trimmed k -means problem is combinatorial and thus difficult to solve exactly. An approximate solution for maximizing (11) is to use the TCLUST algorithm. Here, we adapt the description of the TCLUST algorithm (given in [47, Sec. 3]) for the α -trimmed k -means problem.

Let $\boldsymbol{\gamma}^{(0)}$ denote some randomized initialization of the vector $\boldsymbol{\gamma}$, and let $\boldsymbol{\gamma}^{(r)T} = (\boldsymbol{\mu}_1^{(r)}, \dots, \boldsymbol{\mu}_k^{(r)})$ denote the r th iterate of the algorithm. At the $(r + 1)$ th iterate of the algorithm, perform the steps:

- 1) Compute $d_i^{(r)} = \max_c D_c(\mathbf{u}_i; \boldsymbol{\gamma}^{(r)})$ and store the set of \mathbf{u}_i with the $\lfloor n(1 - \alpha) \rfloor$ largest $d_i^{(r)}$ values in the set $\mathbb{H}^{(r+1)}$.
- 2) Split $\mathbb{H}^{(r+1)}$ into sets $\mathbb{H}_1^{(r+1)}, \dots, \mathbb{H}_k^{(r+1)}$, where

$$\mathbb{H}_c^{(r+1)} = \left\{ \mathbf{u}_i \in \mathbb{H}^{(r+1)} : D_c(\mathbf{u}_i; \boldsymbol{\gamma}^{(r)}) = d_i^{(r)} \right\}.$$

3) For each $c = 1, \dots, k$, compute the updates

$$\mu_c^{(r+1)} = \left| \mathbb{H}_c^{(r+1)} \right|^{-1} \sum_{i=1}^n \mathbf{u}_i \mathbb{I} \left\{ \mathbf{u}_i \in \mathbb{H}_c^{(r+1)} \right\}.$$

Here $\lfloor x \rfloor$ is the floor of $x \in \mathbb{R}$. Steps 1–3 are repeated until some convergence criterion is met upon which the final iterate is declared the α -trimmed k -means estimator $\tilde{\gamma}_n$. In this article, we follow the default option of the *tkmeans* function and terminate the algorithm after 20 iterations.

Theoretical Considerations

The consistency of the α -trimmed k -means estimator $\tilde{\gamma}_n$ is established for the case of independent and identically distributed random variables in [47]; see also [30]. Due to the application to CI, we require that the α -trimmed k -means estimator be consistent for data with some dependence structure. Let \preceq denote the lexicographic order relation in \mathbb{R}^d (cf. [58, Example 2.1.10 (3)]), and let $\Gamma = \{\gamma \in \mathbb{R}^{kd} : \mu_1 \preceq \dots \preceq \mu_k\}$. Further, let $l(\mathbf{u}; \gamma) = \sum_{c=1}^k \delta_c(\mathbf{u}; \gamma) \log D_c(\mathbf{u}; \gamma)$. The following proposition establishes one such result.

Proposition 8. *Let $\mathbf{U}_1, \dots, \mathbf{U}_n \in \mathbb{R}^d$ be an ID and stationary ergodic (or strong-mixing) random sample from a continuous distribution and assume that $|\text{cov}(\mathbf{U}_1)| < \infty$. Assume that there exists some $\gamma_0 \in \Gamma$ such that for every open subset $\mathcal{G} \subset \Gamma$, $\gamma_0 \in \mathcal{G}$ implies that $\mathbb{E}l(\mathbf{U}_1; \gamma_0) > \sup_{\gamma \in \Gamma \setminus \mathcal{G}} \mathbb{E}l(\mathbf{U}_1; \gamma)$. If $\tilde{\gamma}_n \in \Gamma$ is a sequence of estimators such that $\tilde{\gamma}_n = \arg \max_{\gamma \in \Gamma} \mathcal{L}_n(\gamma)$, then $\tilde{\gamma}_n \xrightarrow{P} \gamma_0$.*

Proof: We invoke the M-estimator theorem of [59, Thm. 2.12], which requires we establish that

$$\sup_{\gamma \in \Gamma} |\mathcal{L}_n(\gamma) - \mathbb{E}l(\mathbf{U}_1; \gamma)| \xrightarrow{P} 0.$$

Under the ID stationary ergodic (or strong-mixing) condition, we can use the ULLN of [37, Thm. 5] by verifying that $\mathbb{E}l(\mathbf{U}_1; \gamma) < \infty$ (which also verifies the use of an appropriate law of large numbers; e.g. [36, Thm. 3.34]). Since $\delta_c(\mathbf{u}_1; \gamma)$ are binary, we require that $\mathbb{E} \log D_c(\mathbf{u}; \gamma)$ be finite for each c . Since $\log D_c(\mathbf{u}; \gamma)$ is nonlinear in \mathbf{u} via only quadratic terms, we have $\mathbb{E} \log D_c(\mathbf{u}; \gamma) < \infty$ and hence $\mathbb{E}l(\mathbf{u}; \gamma) < \infty$, under the assumption that $|\text{cov}(\mathbf{U}_1)| < \infty$. ■

Remark 9. The constraint of the parameter space to Γ is require to break the symmetries of potential solutions. Without the restriction, the assumption that there exists some $\gamma_0 \in \Gamma$ such that for every open subset $\mathcal{G} \subset \Gamma$, $\gamma_0 \in \mathcal{G}$ implies that $\mathbb{E}l(\mathbf{U}_1; \gamma_0) > \sup_{\gamma \in \Gamma \setminus \mathcal{G}} \mathbb{E}l(\mathbf{U}_1; \gamma)$ cannot hold, since there will always be a $\gamma \in \Gamma \setminus \mathcal{G}$ such that $\mathbb{E}l(\mathbf{U}_1; \gamma_0) = \mathbb{E}l(\mathbf{U}_1; \gamma)$. The lexicographic ordering restriction is similar to that which is used by [60], in order to break symmetries in mixture of experts models.

Remark 10. As with Proposition 4, an M -dependence assumption can be used to enforce strong-mixing; see Remark 5. However, similarly to Remark 6, M -dependence may be difficult to establish if all dimensions of the calcium imaging is allowed to extend infinitely. An alternative theorem based on the results of [40] can be used instead in such a case.

REFERENCES

- [1] R. D. Burgoyne, “Neuronal calcium sensor proteins: generating diversity in neuronal Ca^{2+} signalling,” *Nature Reviews Neuroscience*, vol. 8, pp. 182–193, 2007.
- [2] W. Denk, R. Yuste, K. Svoboda, and D. W. Tank, “Imaging calcium dynamics in dendritic spines,” *Current Opinion in Neurobiology*, vol. 6, pp. 372–378, 1996.
- [3] R. Yasuda, E. A. Nimchinsky, V. Scheuss, T. A. Pologruto, T. G. Oertner, B. L. Sabatini, and K. Svoboda, “Imaging calcium concentration dynamics in small neuronal compartments,” *Science’s STKE*, vol. 2004, no. 219, 2004.
- [4] T.-W. Chen, T. J. Wardill, Y. Sun, S. R. Pulver, S. L. Renninger, A. Baohuan, E. R. Schreiter, R. A. Kerr, M. B. Orger, V. Jayaraman, and L. L. Looger, “Ultrasensitive fluorescent proteins for imaging neuronal activity,” *Nature*, vol. 499, pp. 295–300, 2013.
- [5] A. M. Stewart, J. F. Ullmann, W. H. Norton, M. O. Parker, C. H. Brennan, R. Gerlai, and A. V. Kalueff, “Molecular psychiatry of zebrafish,” *Molecular Psychiatry*, vol. 20, pp. 2–17, 2015.
- [6] J. F. Ullmann, F. Calamante, S. P. Collin, D. C. Reutens, and N. D. Kurniawan, “Enhanced characterization of the zebrafish brain as revealed by super-resolution track-density imaging,” *Brain Structure and Function*, vol. 220, pp. 457–468, 2015.
- [7] T. R. Thiele, J. C. Donovan, and H. Baier, “Descending control of swim posture by a midbrain nucleus in zebrafish,” *Neuron*, vol. 83, pp. 679–691, 2014.
- [8] A. W. Thompson, G. C. Vanwallendael, L. A. Heap, and E. K. Scott, “Functional profiles of visual-, auditory-, and water flow-responsive neurons in the zebrafish tectum,” *Current Biology*, vol. 26, pp. 743–754, 2016.
- [9] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller, “Whole-brain functional imaging at cell resolution using light-sheet microscopy,” *Nature Methods*, vol. 10, pp. 413–420, 2013.
- [10] G. D. Marquart, K. M. Tabor, M. Brown, J. L. Strykowski, G. K. Varshney, M. C. LaFave, T. Mueller, S. M. Burgess, S. I. Higashijima, and H. A. Burgess, “A 3D search database of transgenic zebrafish Gal4 and Cre lines for functional neuroanatomy studies,” *Frontiers in Neural Circuits*, vol. 9, 2015.
- [11] M. B. Bouchard, V. Voleti, C. S. Mendes, C. Lacefield, W. B. Grueber, R. S. Mann, R. M. Bruno, and E. M. C. Hillman, “Swept confocally-aligned planar excitation (SCAPE) microscopy for high-speed volumetric imaging of behaving organisms,” *Nature Photonics*, vol. 9, pp. 113–119, 2015.
- [12] M. Chen, S. Mao, and Y. Liu, “Big Data: a survey,” *Mobile Networks and Applications*, vol. 19, pp. 171–209, 2014.
- [13] T. W. Liao, “Clustering of time series data—a survey,” *Pattern Recognition*, vol. 38, pp. 1857–1874, 2005.
- [14] P. Esling and C. Agon, “Time-series data mining,” *ACM Computing Surveys*, vol. 45, pp. 12:1–12:34, 2012.
- [15] J. O. Ramsay and B. W. Silverman, *Functional Data Analysis*. New York: Springer, 2005.
- [16] S. Gaffney and P. Smyth, “Trajectory clustering with mixtures of regression models,” in *Proceedings of the fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1999.
- [17] G. M. James and C. A. Sugar, “Clustering for sparsely sampled functional data,” *Journal of the American Statistical Association*, vol. 98, pp. 397–408, 2003.
- [18] G. Celeux, O. Martin, and C. Lavergne, “Mixture of linear mixed models for clustering gene expression profiles from repeated microarray experiments,” *Statistical Modelling*, vol. 5, pp. 243–267, 2005.
- [19] S. K. Ng, G. J. McLachlan, K. W. L. Ben-Tovim, and S. W. Ng, “A mixture model with random-effects components for clustering correlated gene-expression profiles,” *Bioinformatics*, vol. 22, pp. 1745–1752, 2006.
- [20] K. Wang, S. K. Ng, and G. J. McLachlan, “Clustering of time-course gene expression profiles using normal mixed models with autoregressive random effects,” *BMC Bioinformatics*, vol. 13, p. 300, 2012.
- [21] S. K. Ng and G. J. McLachlan, “Mixture of random effects models for clustering multilevel growth trajectories,” *Computational Statistics and Data Analysis*, vol. 71, pp. 43–51, 2014.
- [22] H. D. Nguyen, G. J. McLachlan, and I. A. Wood, “Mixture of spatial spline regressions for clustering and classification,” *Computational Statistics and Data Analysis*, vol. 93, pp. 76–85, 2016.
- [23] Y. Xiong and D.-Y. Yeung, “Time series clustering with ARMA mixtures,” *Pattern Recognition*, vol. 37, pp. 1675–1689, 2004.
- [24] H. D. Nguyen, G. J. McLachlan, J. F. P. Ullmann, and A. L. Janke, “Spatial clustering of time-series via mixture of autoregressions models and Markov Random Fields,” *Statistica Neerlandica*, vol. In Press, 2016.
- [25] H. D. Nguyen, G. J. McLachlan, P. Orban, P. Bellec, and A. L. Janke, “Maximum pseudolikelihood estimation for model-based clustering of time series data,” *arXiv:1602.08787*, 2016.

- [26] H. D. Nguyen, G. J. McLachlan, J. F. P. Ullmann, and A. L. Janke, "Model-based clustering of functional data via Gaussian mixture models," *arXiv:1608.05481*, 2016.
- [27] G. J. McLachlan and D. Peel, *Finite Mixture Models*. New York: Wiley, 2000.
- [28] G. Schwarz, "Estimating the dimensions of a model," *Annals of Statistics*, vol. 6, pp. 461–464, 1978.
- [29] J.-P. Baudry, C. Maugis, and B. Michel, "Slope heuristic: overview and implementation," *Statistics and Computing*, vol. 22, pp. 455–470, 2012.
- [30] J. A. Cuesta-Albertos, A. Gordaliza, and C. Matran, "Trimmed k-means: an attempt to robustify quantizers," *Annals of Statistics*, vol. 25, pp. 553–576, 1997.
- [31] C. Abraham, P. A. Cornillon, E. Matzner-Lober, and N. Molinari, "Unsupervised curve clustering using B-splines," *Scandinavian Journal of Statistics*, vol. 30, pp. 581–595, 2003.
- [32] J. Peng and H.-G. Muller, "Distance-based clustering of sparsely observed stochastic processes, with applications to online auctions," *Annals of Applied Statistics*, vol. 2, pp. 1056–1077, 2008.
- [33] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society Series B*, vol. 39, pp. 1–38, 1977.
- [34] C. de Boor, *A Practical Guide to Splines*. New York: Springer, 2001.
- [35] T. Amemiya, *Advanced Econometrics*. Cambridge: Harvard University Press, 1985.
- [36] H. White, *Asymptotic Theory For Econometricians*. San Diego: Academic Press, 2001.
- [37] D. W. K. Andrews, "Generic uniform convergence," *Econometric Theory*, vol. 8, pp. 241–257, 1992.
- [38] R. C. Bradley, "Basic properties of strong mixing conditions. A survey and some open questions," *Probability Surveys*, vol. 2, pp. 107–144, 2005.
- [39] —, "A caution on mixing conditions for random fields," *Statistics and Probability Letters*, vol. 8, pp. 489–491, 1989.
- [40] N. Jenish and I. R. Prucha, "Central limit theorems and uniform laws of large number for arrays of random fields," *Journal of Econometrics*, 2009.
- [41] J. A. Hartigan and M. A. Wong, "Algorithm AS 136: A k-means clustering algorithm," *Journal of the Royal Statistical Society Series C*, vol. 28, pp. 100–108, 1979.
- [42] G. J. McLachlan, *Discriminant Analysis And Statistical Pattern Recognition*. New York: Wiley, 1992.
- [43] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*. Cambridge: Cambridge University Press, 2007.
- [44] P. Massart, *Concentration Inequalities and Model Selection*. New York: Springer, 2007.
- [45] L. Birge and P. Massart, "Minimal penalties for Gaussian model selection," *Probability Theory and Related Fields*, vol. 138, pp. 33–73, 2007.
- [46] J.-P. Baudry, C. Maugis, and B. Michel, "Slope heuristic: overview and implementation," INRIA, Tech. Rep. hal-00461639, 2010.
- [47] L. A. Garcia-Escudero, A. Gordaliza, C. Matran, and A. Mayo-Isar, "A general trimming approach to robust cluster analysis," *Annals of Statistics*, 2008.
- [48] N. Vladimirov, Y. Mu, T. Kawashima, D. V. Bennett, C. T. Yang, L. L. Looger, P. J. Keller, J. Freeman, and M. B. Ahrens, "Light-sheet functional imaging in fictively behaving zebrafish," *Nature Methods*, vol. 11, pp. 883–884, 2014.
- [49] R Core Team, *R: a language and environment for statistical computing*, R Foundation for Statistical Computing, 2013.
- [50] R. D. Vincent, A. Janke, J. G. Sled, L. Baghdadi, P. Neelin, and A. C. Evans, "MINC 2.0: a modality independent format for multidimensional medical images," in *10th Annual Meeting of the Organization for Human Brain Mapping*, 2003.
- [51] C. Bordier, M. Dojat, and P. L. de Micheaux, "Temporal and spatial independent component analysis for fMRI data sets embedded in the AnalyzeFMRI R package," *Journal of Statistical Software*, vol. 44, pp. 1–24, 2011.
- [52] M. Enea, *speedglm: fitting linear and generalized linear models to large data sets*, 2015. [Online]. Available: <https://CRAN.R-project.org/package=speedglm>
- [53] J. O. Ramsay, G. Hooker, and S. Graves, *Functional Data Analysis with R and MATLAB*. New York: Springer, 2009.
- [54] H. Fritz, L. A. Garcia-Escudero, and A. Mayo-Isar, "TCLUST: an R package for a trimming approach to cluster analysis," *Journal of Statistical Software*, vol. 47, pp. 1–26, 2012.
- [55] S. Arlot, V. Brault, J.-P. Baudry, C. Maugis, and B. Michel, *capushe: CALibrating Penalties Using Slope HEuristics*, 2015. [Online]. Available: <https://CRAN.R-project.org/package=capushe>
- [56] D. Eddelbuettel, *Seamless R and C++ Integration with Rcpp*. New York: Springer, 2013.

- [57] H. D. Nguyen, G. J. McLachlan, N. Cherbuin, and A. L. Janke, “False discovery rate control in magnetic resonance imaging studies via Markov random fields,” *IEEE Transactions on Medical Imaging*, vol. 33, pp. 1735–1748, 2014.
- [58] A. A. Khan, C. Tammer, and C. Zalinescu, *Set-valued Optimization*. New York: Springer, 2015.
- [59] M. R. Kosorok, *Introduction to Empirical Processes and Semiparametric Inference*. New York: Springer, 2008.
- [60] W. Jiang and M. A. Tanner, “On the identifiability of mixture-of-experts,” *Neural Networks*, vol. 12, pp. 1253–1258, 1999.